arXiv:1704.07078v1 [cs.SI] 24 Apr 2017

Rethinking (k, ℓ)-anonymity in social graphs: (k, ℓ)-adjacency anonymity and (k, ℓ)-(adjacency) anonymous transformations

S. Mauw[1], Y. Ramírez-Cruz[2] and R. Trujillo-Rasua[2]

[1]CSC, [2]SnT

University of Luxembourg

6, av. de la Fonte, L-4364 Esch-sur-Alzette, Luxembourg

{sjouke.mauw, yunior.ramirez, rolando.trujillo}@uni.lu

April 25, 2017

Abstract

This paper treats the privacy-preserving publication of social graphs in the presence of active adversaries, that is, adversaries with the ability to introduce sybil nodes in the graph prior to publication and leverage them to create unique fingerprints for a set of victim nodes and re-identify them after publication. Stemming from the notion of (k, ℓ)-anonymity, we introduce (k, ℓ)-anonymous transformations, characterising graph perturbation methods that ensure protection from active adversaries levaraging up to ℓ sybil nodes. Additionally, we introduce a new privacy property: (k, ℓ)-adjacency anonymity, which relaxes the assumption made by (k, ℓ)-anonymity that adversaries can control all distances between sybil nodes and the rest of the nodes in the graph. The new privacy property is in turn the basis for a new type of graph perturbation: (k, ℓ)-adjacency anonymous transformations. We propose algorithms for obtaining $(k, 1)$-adjacency anonymous transformations for arbitrary values of k, as well as $(2, \ell)$-adjacency anonymous transformations for small values of ℓ.

Keywords: social graphs, privacy-preserving publication, active adversaries, perturbation methods.

1 Introduction

Online social networks (OSNs) have become the most successful application of our time. Nearly two billion persons[1] regularly use some OSN to interact with friends and relatives, share information, get news, entertainment, etc. As a result of this, massive

[1]Source: statista.com, consulted on April 21st, 2017.

amounts of information about human behaviour, personal relationships, consumption patters, personal preferences, and more, are generated everyday. An important part of this information is encoded in the form of social graphs. In a social graph, every vertex corresponds to a person (a user of the OSN), whereas edges represent relations between individuals. Rich personal information, such as name, address, etc. is usually associated to vertices as attributes. Edges can also have associated attributes, which may encode, for instance, the nature of the relation (friendship, laboral, family), the date and place where it was established, etc.

This massive amount of information is enormously valuable. OSNs themselves analyse this data in order to determine the advertisement they show to their users, suggest new potential friends, filter out the information that they consider not to be interesting to the user, etc. As holders of the information, the OSN can effectively access the totality of the available data, as authorised by the users when they upload their information. However, third parties, such as social scientists, market researchers, public institutions and private companies, are also interested in accessing and analysing a part of this information for conducting population studies, assessing the effect of communication campaigns, surveying public opinion, and many other purposes. In order to enable these studies, it is necessary that the OSN administrators release samples of their social graphs. However, despite the usefulness of the studies that can be conducted on the released data, the sensitive nature of a part of the information encoded in social graphs, e.g. political or religious affiliation, arises serious privacy concerns.

A *naïve* approach to protect the privacy of users in publishing social graphs is to remove all personally identifiable information from the released graph. However, as shown in [6], even a graph with no identifying attributes can leak sensitive information, since some structural properties (the degree of vertices, their neighbourhoods, etc.) can be unique to certain users. A *re-identification attack* seeks to leverage some knowledge about a set of users, the victims, to re-identify them after the graph is released. For example, an adversary who knows the number of friends of all the victim vertices can later re-identify them in the released graph if every value happens to be unique, even if all vertex and edge attributes have been removed. Once a set of users is re-identified, the attacker can learn sensitive information, such as the existence of relations between two users or the (co-)affiliation of some of them to a community.

According to the means by which adversaries obtain the knowledge used to re-identify the victims, they are classified as *passive* or *active* [1]. Passive adversaries rely on information obtainable from publicly available sources, such as other OSNs, but do not attempt to purposely alter the structure of the network. On the other hand, active adversaries enroll sybil nodes in the network and try to force the creation of structural patterns that allow them to later re-identify the victims. Several active attacks were described in [1, 7, 8]. In every attack, the adversary inserts a small number of sybil nodes, and then creates unique connection patters, referred to as *fingerprints*, between sybil nodes and the victims. Additionally, the connections between pairs of sybil nodes are established in such a way that the subgraph induced by them is not isomorphic (with high probability) to any other subgraph. Once the graph is published, the uniquely identifiable set of sybil nodes is retrived, and the unique fingerprints allow

to re-identify the victims.

The notion of (k, ℓ)-anonymity was introduced in [11] as a measure of the resistance of a social graph to active attacks. Informally, a (k, ℓ)-anonymous graph ensures that an adversary with the ability to insert up to ℓ sybil nodes in the network, cannot use the distances from these sybil nodes to other vertices to uniquely identify any vertex. This guarantee comes from the fact that each vertex is ensured to be undistinguishable from at least other $k - 1$ vertices according to the so-called metric representation with respect to every vertex subset of size at most ℓ. Then, a family of methods that transform a $(1, 1)$-anonymous graph G (which is the least private type of graphs) into a graph G' that satsifies (k, ℓ)-anonymity for $k > 1$ or $\ell > 1$, was proposed in [5, 4].

In this paper, we re-visit the notion of (k, ℓ)-anonymity. We focus on two assumptions encoded in (k, ℓ)-anonymity: treating every vertex subset of size up to ℓ as a potential set of sybil nodes, and assuming that the adversary is able to control the distances between the set of sybil nodes and every other vertex in the graph, which is not realistic. As a result, we first propose the notion of (k, ℓ)-anonymous transformations, which ensure the same level of protection that would be achieved by enforcing (k, ℓ)-anonymity while performing less modifications in the graph. Then, we introduce a new privacy property, (k, ℓ)-adjacency anonymity, which relaxes the assumption on the distances that the adversary is able to control. Finally, these two ideas are combined in the notion of (k, ℓ)-adjacency anonymous transformation, and we propose two methods, based on edge additions and removals, for performing $(k, 1)$- and $(2, \ell)$-adjacency anonymous transformations.

The remainder of this paper is structured as follows. In Section 2 we discuss our new adversary model, introducing (k, ℓ)-anonymous transformations, (k, ℓ)-adjacency anonymity, and (k, ℓ)-adjacency anonymous transformations. Section 3 introduces the algorithm for obtaining $(k, 1)$-adjacency anonymous transformations, whereas Section 4 introduces the algorithm for obtaining $(2, \ell)$-adjacency anonymous transformations. Finally, we discuss our results and possible directions for future work in Section 5. Before proceeding, we will introduce some notation that will be used throughout the paper. We will use the notation $u \sim_G v$ for two vertices $u, v \in V(G)$ that are adjacent in G, *i.e.* $(u, v) \in E(G)$. The *open neighbourhood* of a vertex $u \in V(G)$, denoted by $N_G(u)$, is the set $N_G(u) = \{v : u \sim_G v\}$, whereas the *closed neighbourhood* of u in G is the set $N_G[u] = \{u\} \cup N_G(u)$. Similarly, for a set $S \subset V(G)$, we define $N_G(S) = \cup_{v \in S} N_G(v) \setminus S$ and $N_G[S] = \cup_{v \in S} N_G[v]$. The *degree* of a vertex u, denoted by $\delta_G(u)$, is its number of neighbours, i.e. $\delta_G(u) = |N_G(u)|$. In a graph G of order n, we will refer to vertices of degree 0, 1 and $n - 1$ as *isolated*, *end-* and *dominant* vertices, respectively. The *distance* between two vertices u and v in G, denoted as $d_G(u, v)$, is the number of edges in a shortest path joining u and v. For a graph $G = (V, E)$ and a subset S of vertices of G, we will denote by $\langle S \rangle_G$ the subgraph of G induced by S, that is $\langle S \rangle_G = (S, E \cap S \times S)$. In the previously defined notations, if there is no ambiguity, we will drop the graph-specific subindices and simply write $u \sim v$, $N(u)$, $\delta(u)$, etc. For a graph G, we define $\delta(G) = \min_{v \in V(G)}\{\delta_G(v)\}$ and $\Delta(G) = \max_{v \in V(G)}\{\delta_G(v)\}$ and, as usual, we will denote by K_n and N_n the complete and empty graphs of order n, respectively.

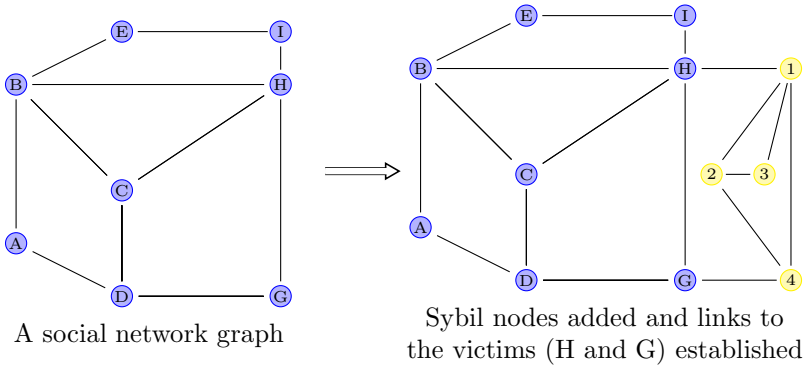

(a) Active attack prior publication.

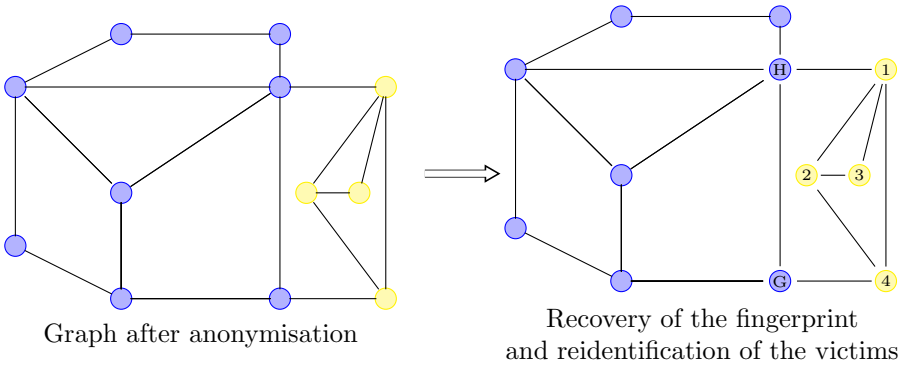

(b) Reidentification after publication.

Figure 1: The four stages of an active attack.

2 Adversary model

An active adversary uses graph properties of a set of sybil nodes to re-identify users in an anonymised social graph. Prior publication of the social network graph, the active attacker adds a set of sybil nodes to the network (e.g. nodes $1, 2, 3$ and 4 in Figure 1a). The sybil nodes establish links between themselves and also with the victims (e.g. users H and G in Figure 1a). After publication of the social network graph without the users' identifiers, the attacker first searches for the subgraph formed by the sybil nodes. Victims connected to the attacker subgraph can be reidentified by using the neighbour relation between sybil nodes and victims. For example, the non-sybil nodes connected to 1 and 4 in Figure 1b, respectively, must be H and G. This allows the adversary to acquire knowledge that was supposed to remain private, such as the existence of a link between users H and G.

From a practical point of view, active attacks require the ability to insert sybil nodes in the social network and remain unnoticed by sybil detection techniques. This is a fairly easy task in today's social networks, as false positives in sybil detection are undesirable and registration to the network should be trivial; social networks understandably favour usability and user experience over sybil detection. From a theoretical point of view, an active attack relies on creating a *unique* attacker sub-

graph. That is to say, the induced subgraph formed by the sybil nodes should have no trivial automorphism and no other subgraph in the network isomorphic to it. For example, assume that the adversary in the attack in Figure 1 could not insert the third node, i.e. the node labelled 3. This makes the subgraph induced by $1, 2$ and 4 isomorphic to the subgraph induced by B, C, and H, which prevents the attacker from correctly retrieving the inserted subgraph.

Backstrom et al. already showed that, despite of the previously mentioned challenges, active attacks can be implemented successfully [1]. They proved that only $\log n$ sybil nodes, where n is the number of vertices of the network, are needed to create an attacker subgraph which is unique with high probability. That makes active attacks particularly dangerous and hard to prevent.

Effectively determining whether a social graph is vulnerable to an active attack is a necessary step towards developing a mitigation strategy against it. For example, the complete graph G satisfies that for every proper subgraph S there exists another subgraph S' that is isomorphic to S. Such property makes an active attack unfeasible in a complete graph. Determining the actual resistance of an arbitrary graph to active attacks is not trivial, though. A first step on this direction was given in [11], where Trujillo-Rasua and Yero introduced the privacy measure (k, ℓ)-anonymity.

Consider a total order $\preceq$ on the vertices of a graph G. Given a set $S \subseteq V(G)$, let $(v_1, v_2, \ldots, v_t)$, where $v_i \in S$ for $i \in \{1, \ldots, t\}$, be the vector composed by the elements of S, in such a way that $v_1 \preceq v_2 \preceq \ldots \preceq v_t$. For the sake of simplicity in our presentation, in what follows we will abuse notation and refer to the *ordered set* $S = \{v_1, v_2, \ldots, v_t\}$. Given an ordered set of sybil nodes $S = (s_1, \ldots, s_t)$ in a graph $G = (V, E)$, Trujillo-Rasua and Yero [11] define the adversary knowledge about a user $u \in V$ as the vector $(d_G(v, s_1), \ldots, d_G(v, s_t))$. This vector is referred to as *metric representation* of u with respect to S, and denoted $r_G(u|S)$ [10, 2].

The metric representation was introduced in [10, 2] as a tool to define the so-called *resolving sets*. A set $S \subset V(G)$ is said to be a resolving set[2] of G if every vertex $u \in V(G)$ has a unique metric representation with respect to S. This property of resolving sets inspired the definition in [11] of an opposed concept, with implications in vertex privacy.

Definition 1 (k-antiresolving set)**.** *Let $G = (V, E)$ be a non-trivial graph. A set $S \subset V$ is a* k-antiresolving set *of G if k is the largest positive integer such that, for every $v \in V(G) \setminus S$, there exist vertices $w_1, w_2, \ldots, w_{k-1} \in V(G) \setminus S$ such that $v, w_1, w_2, \ldots, w_{k-1}$ are pairwise different and*

$$r_G(v \mid S) = r_G(w_1 \mid S) = r_G(w_2 \mid S) = \ldots = r_G(w_{k-1} \mid S).$$

The minimum cardinality of a k-antiresolving set of a graph G is called the k-metric antidimension of G. These concepts were used to quantify the privacy of a social graph in the presence of active attackers as follows.

[2]Multiple terminologies have been used for resolving sets in the literature. The term *resolving set* was introduced by Slater in [10], whereas Harary and Melter [2] introduced the terminology *locating set* and Sebő and Tannier [9] preferred the term *metric generator*.

Definition 2 ((k,ℓ)-anonymity)**.** *A graph G is said to satisfy* (k,ℓ)-anonymity *if k is the smallest positive integer such that the k-metric antidimension of G is smaller than or equal to ℓ.*

From a privacy perspective, if a graph satisfies (k,ℓ)-anonymity, an attacker with the capacity to enrol, and successfully retrieve, up to ℓ sybil nodes in the graph would still be incapable of distinguishing any vertex from at least other $k-1$ vertices. Taking back again the example of the complete graph K_n, it is easy to prove that K_n satisfies $(n-l,l)$-anonymity. It is worth noticing that $k=n-l$ corresponds to the maximum value possible for k in (k,ℓ)-anonymity given $\ell=l$.

Certainly, a graph satisfying (k,ℓ)-anonymity for $k>1$ effectively resists active attacks when performed by at most ℓ sybil nodes. However, event the simplest of the privacy goals, namely transforming a $(1,1)$-anonymous graph into a (k,ℓ)-anonymous graph G' with either $k>1$ or $\ell>1$, has not been accomplished without significant information loss [5]. Our observation is that (k,ℓ)-anonymity, although suitable to quantify resistance against active attacks, cannot be applied straightforwardly to privacy-preserving transformation of social graphs.

2.1 Revisiting (k,ℓ)-anonymity

(k,ℓ)-anonymity quantifies over all subsets of vertices of size at most ℓ. Therefore, a transformation from an original graph G to an anonymized graph G' satisfying, for example $(2,\ell)$-anonymity, must ensure that every subset of vertices S in G' with $|S|\leq\ell$ is a k'-antiresolving set where $k'\geq 2$, regardless of whether S was indeed a 1-antiresolving set in G or not. In effect, assuming that the set of attacker nodes S is already a 2-antiresolving set in the original graph G, it is harmless to publish (with respect to $(2,|S|)$-anonymity) a transformation of G where S is a 1-antiresolving set. Consequently, when aiming at $(2,\ell)$-anonymity, a transformation method should only be concerned about those 1-antiresolving sets with size at most ℓ. We formalise this concept as follows.

Definition 3 ((k,ℓ)-anonymous transformation)**.** *A pair (G_1,G_2) is a* (k,ℓ)-anonymous transformation *if for every subset $S\subseteq V(G_1)\cap V(G_2)$ with $|S|\leq\ell$, S being a k_1-antiresolving set in G_1, and S being k_2-antiresolving set in G_2, it holds that $k_1<k \implies k_2\geq k$.*

Notice that, in particular, if a graph G satisfies (k,ℓ)-anonymity, then every pair (G_0,G), where G_0 is an arbitrary graph, is a (k,ℓ)-anonymous transformation. The converse is not true, as exemplified in Figure 2.

2.2 The adversary knowledge

Privacy measures based on k-anonymity are defined based on a concrete definition of the *adversary knowledge*. In (k,ℓ)-anonymity, an adversary is a set of sybil nodes $S\subseteq V(G)$ within a network G. The knowledge of such adversary about a user $u\in V(G)-S$ is considered to be the metric representation $r_G(u|S)$. That is to say,

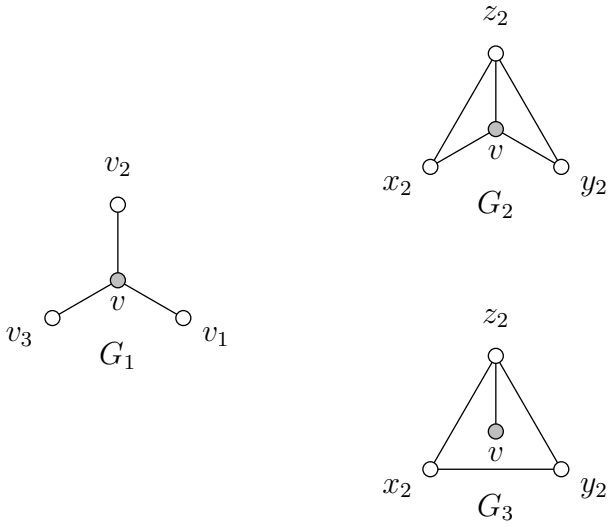

Figure 2: Two $(2,1)$-anonymous transformations (G_1, G_2) and (G_1, G_3). The graph G_2 satisfies $(2,1)$-anonymity, whereas G_3 does not.

the adversary is capable of determining the distance from every attacker node to any other node in the network. This is a strong assumption, yet it can be justified by the necessity of not underestimating the adversary capabilities.

In this article we relax the assumption on the adversary knowledge made in [11]. Our decision is based on the fact that all active attacks proposed so far [1, 7, 8] rely on the neighbour relation between the attacker nodes and the victims. It is indeed unrealistic to expect the adversary to rely on arbitrary distances, since that would imply knowing the entire adjacency matrix and, especially, having the capability to influence whether a relation is established, or not, between any pair of users of the network.

In a manner analogous to the definition of antiresolving sets, we use standard concepts from Graph theory to represent an adversary that only has knowledge about its neighbours. The concept is known as *adjacency representation*, introduced by Jannesari and Omoomi [3] and defined as follows.

Definition 4 (Adjacency representation)**.** *Given a graph $G = (V, E)$, an ordered set $S = \{s_1, \ldots, s_t\} \subset V$, and a vertex $u \in V$, the* adjacency metric representation *of u with respect S is the vector $a_G(r|S) = (a_G(s_1, v), \ldots, a_G(s_t, v))$ where $a_G : V(G) \times V(G) \to \mathbb{N}$ is defined by:*

$$a_G(u,v) = \begin{cases} 0 & \text{if } u = v \\ 1 & \text{if } u \sim_G v \\ 2 & \text{otherwise} \end{cases} \tag{1}$$

Note that $a_G(u,v) = \min\{2, d_G(u,v)\}$ for every $u, v \in V(G)$. Now, we will adapt the notion of k-antiresolving sets in order to account for the new type of adversary.

Definition 5 (k-adjacency antiresolving set)**.** *Let $G = (V, E)$ be a non-trivial graph. A set $S \subset V$ is a* k-adjacency antiresolving set *of G if k is the largest positive integer such that, for every $v \in V(G) \setminus S$, there exist vertices $w_1, w_2, \ldots, w_{k-1}$ such that $v, w_1, w_2, \ldots, w_{k-1}$ are pairwise different and*

$$a_G(v \mid S) = a_G(w_1 \mid S) = a_G(w_2 \mid S) = \ldots = a_G(w_{k-1} \mid S).$$

To illustrate the difference between k-adjacency antiresolving sets and k-antiresolving sets, consider the graph G depicted in Figure 3. The set $\{v\}$ is a 2-antiresolving set of G, as z_1 and z_2 satisfy $d(v, z_1) = d(v, z_2) = 3$, whereas $d(v, x_1) = d(v, x_2) = d(v, x_3) = d(v, x_4) = 1$ and $d(v, y_1) = d(v, y_2) = d(v, y_3) = d(v, y_4) = 2$. On the other hand, we have that $a_G(x_1, \{v\}) = a_G(x_2, \{v\}) = a_G(x_3, \{v\}) = a_G(x_4, \{v\}) = (1)$, while $a_G(y_1, \{v\}) = a_G(y_2, \{v\}) = a_G(y_3, \{v\}) = a_G(y_4, \{v\}) = a_G(z_1, \{v\}) = a_G(z_2, \{v\}) = (2)$, so $\{v\}$ is a 4-adjacency antiresolving set of G.

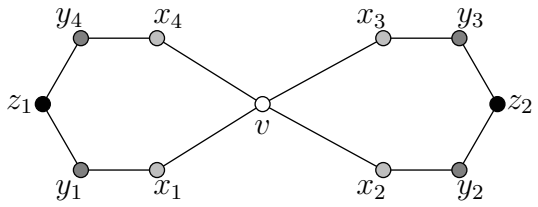

Figure 3: In this graph, the set $\{v\}$ is a 2-antiresolving set and a 4-adjacency antiresolving set.

For a graph $G = (V, E)$ and a set $S \subset V$, let $R_{G,S}$ be the equivalence relation such that two vertices u and v satisfy $u \; R_{G,S} \; v$ if and only if $u, v \in V \setminus S$ and $a_G(u \mid S) = a_G(v \mid S)$. Moreover, we will use the notation $\mathcal{A}_{G,S}$ for the set of equivalence classes induced in $V \setminus S$ by the relation $R_{G,S}$. It is simple to see that S is a $(\min_{A \in \mathcal{A}_{G,S}}\{|A|\})$-adjacency antiresolving set of G.

2.3 Problem statement

We will first enunciate the notions of (k, ℓ)-adjacency anonymity and (k, ℓ)-adjacency anonymous transformation, which restrict the original definitions of (k, ℓ)-anonymity and (k, ℓ)-anonymous transformation, to account for adversaries whose knowledge consists of the adjacency representations of their victims.

Definition 6 (k-adjacency antidimension)**.** *The k-adjacency antidimension of a graph G is the minimum cardinality of a k-adjacency antiresolving set of G.*

Definition 7 ((k, ℓ)-adjacency anonymity)**.** *A graph G satisfies (k, ℓ)-adjacency anonymity if k is the smallest positive integer such that the k-adjacency antidimension of G is smaller than or equal to ℓ.*

According to Definition 7, if a graph G satisfies (k, ℓ)-adjacency anonymity, then for every $S \subset V(G)$ of size at most ℓ and every $u \in V(G) \setminus S$ there exist $v_1, \ldots, v_{k-1} \in$

$V(G) \setminus (S \cup \{u\})$ such that $u, v_1, \ldots, v_{k-1}$ are pairwise different and $a_G(u \mid S) = a_G(v_1 \mid S) = \ldots = a_G(v_{k-1} \mid S)$, so the probability of S being able to re-identify v is at most $1/k$.

It is simple to see that the complete graph K_n and the empty graph $N_n = (V, \emptyset)$ satisfy $(n-\ell, \ell)$-adjacency anonymity for every $\ell \in \{1, n-1\}$, because for every $S \subset V$ the adjacency representation of every other vertex with respect to S is either $(1, 1, \ldots, 1)$ or $(2, 2, \ldots, 2)$, respectively. In the next sections we will introduce results characterising the graphs that satisfy (k, ℓ)-adjacency anonymity for other values of k and ℓ.

Definition 8 ((k, ℓ)-adjacency anonymous transformation)**.** *A pair (G_1, G_2) is a* (k, ℓ)-adjacency anonymous transformation *if for every subset $S \subseteq V(G_1) \cap V(G_2)$ with $|S| \leq \ell$, S being a k_1-adjacency antiresolving set in G_1, and S being k_2-adjacency antiresolving set in G_2, it holds that $k_1 < k \implies k_2 \geq k$.*

In a manner analogous to (k, ℓ)-anonymous transformations, we have that if a graph G satisfies (k, ℓ)-adjacency anonymity, then every pair (G_0, G), where G_0 is an arbitrary graph, is a (k, ℓ)-adjacency anonymous transformation.

Definition 9 (**Problem statement**)**.** *Let* $\text{loss}(G, G')$ *be a cost function providing the information loss incurred by the graph transformation from G to G'. Given a graph G, and natural numbers k and ℓ, find G' such that (G, G') is a (k, ℓ)-adjacency anonymous transformation and* $\text{loss}(G, G')$ *is minimum.*

3 $(k, 1)$-adjacency anonymous transformations

Consider a $(k_0, 1)$-adjacency anonymous graph G of order n. In order to increase the resistance of G to active attackers leveraging one sybil node, our interest is to propose $(k, 1)$-adjacency transformations of the form (G, G') where $k > k_0$. The next result allows us to assess the values of k that may be of interest.

Proposition 10. *Let G be a non-complete, non-empty graph of order n satisfying $(k, 1)$-adjacency anonymity. Then, $k \leq \left\lfloor \frac{n-1}{2} \right\rfloor$.*

Proof. Let $G = (V, E)$ be a non-complete, non-empty graph of order n satisfying $(k, 1)$-adjacency anonymity. Suppose, for the purpose of contradiction, that $k > \left\lfloor \frac{n-1}{2} \right\rfloor$. Let $v \in V$ be a vertex of G satisfying $\delta(v) \notin \{0, n-1\}$. The existence of such a vertex is guaranteed by the fact that the graph is not complete nor empty. We have that $\mathcal{A}_{G,\{v\}} = \{N_G(v), V \setminus N_G[v]\}$. If $\delta(v) \leq \left\lfloor \frac{n-1}{2} \right\rfloor$, then $\{v\}$ is a k'-adjacency antiresolving set of G with $k' < k$, which is a contradiction. On the other hand, if $\delta(v) > \left\lfloor \frac{n-1}{2} \right\rfloor$, then $|V \setminus N_G[v]| \leq \left\lfloor \frac{n-1}{2} \right\rfloor$, which again means that $\{v\}$ is a k'-adjacency antiresolving set of G with $k' < k$, a contradiction. Therefore, we have that $k \leq \left\lfloor \frac{n-1}{2} \right\rfloor$. □

According to Proposition 10, in order to enforce $(k, 1)$-adjacency anonymity on G for some $k > \left\lfloor \frac{n-1}{2} \right\rfloor$, it is necessary to transform G into a complete or empty

graph, which lacks interest for us because such a graph would be completely useless for analysis. Thus, we will focus on the values of k in the interval $\left[k_0+1, \left\lfloor \frac{n-1}{2} \right\rfloor\right]$.

The following results show the relations between the minimum and maximum degrees of a graph and its resistance against active attackers leveraging one sybil node. We first introduce some additional notation. For a graph $G = (V, E)$, let $I_G = \{v \in V : \ \delta(v) = 0\}$ be the set of isolated vertices and let $D_G = \{v \in V : \delta(v) = n-1\}$ be the set of dominant vertices. Clearly, either $I_G = D_G = \emptyset$, or $I_G = \emptyset \wedge D_G \neq \emptyset$, or $I_G \neq \emptyset \wedge D_G = \emptyset$. With these definitions in mind, we give the following three results.

Proposition 11. *Every non-complete graph G n such that $I_G = \emptyset$ satisfies $(k,1)$-adjacency anonymity with $k = \min\{\delta(G), n - \Delta(G) - 1\}$.*

Proof. Let $G = (V, E)$ be a graph of order n without isolated or dominant vertices. Consider a vertex $v \in V$. Clearly, $\mathcal{A}_{G,\{v\}} = \{N_G(v), V \setminus N_G[v]\}$, so the set $\{v\}$ is a $(\min\{\delta(v), n-\delta(v)-1\})$-adjacency antiresolving set of G. In consequence, the graph G satisfies $(k,1)$-adjacency anonymity with

$$\begin{aligned} k &= \min_{v\in V}\{\min\{\delta(v), n-\delta(v)-1\}\} \\ &= \min\left\{\min_{v\in V}\{\delta(v), n-\delta(v)-1\}\right\} \\ &= \min\left\{\min_{v\in V}\{\delta(v)\}, \min_{v\in V}\{n-\delta(v)-1\}\right\} \\ &= \min\{\delta(G), n-\Delta(G)-1\} \end{aligned}$$

□

Proposition 12. *Let $G = (V, E)$ be a non-complete graph of order n such that $D_G \neq \emptyset$ and let $S = V \setminus D_G$. Then, G satisfies $(k,1)$-adjacency anonymity with*

$$k = \min\left\{\delta(G), |S| - \Delta(\langle S\rangle_G) - 1\right\}.$$

Proof. We follow a reasoning analogous to that of the proof of Proposition 11. First, consider a vertex $v \in D_G$. We have that $\mathcal{A}_{G,\{v\}} = \{V \setminus \{v\}\}$. Now, consider a vertex $v \in S$. In this case, $\mathcal{A}_{G,\{v\}} = \{N_G(v), V \setminus N_G[v]\} = \{N_G(v), S \setminus N_{\langle S\rangle_G}[v]\}$. In consequence, we have that G satisfies $(k,1)$-adjacency anonymity with

$$\begin{aligned} k &= \min\left\{n-1, \min_{v\in S}\left\{\min\left\{\delta_G(v), |S| - \delta_{\langle S\rangle_G}(v) - 1\right\}\right\}\right\} \\ &= \min_{v\in S}\left\{\min\left\{\delta_G(v), |S| - \delta_{\langle S\rangle_G}(v) - 1\right\}\right\} \\ &= \min\left\{\min_{v\in S}\left\{\delta_G(v), |S| - \delta_{\langle S\rangle_G}(v) - 1\right\}\right\} \\ &= \min\left\{\min_{v\in S}\{\delta_G(v)\}, \min_{v\in S}\left\{|S| - \delta_{\langle S\rangle_G}(v) - 1\right\}\right\} \\ &= \min\{\delta(G), |S| - \Delta(\langle S\rangle_G) - 1\} \end{aligned}$$

□

Proposition 13. *Let $G = (V, E)$ be a non-empty graph of order n such that $I_G \neq \emptyset$ and let $S = V \setminus I_G$. Then, G satisfies $(k, 1)$-adjacency anonymity with*

$$k = \min \left\{ \delta(\langle S \rangle_G), n - \Delta(G) - 1 \right\}.$$

Proof. We follow a reasoning analogous to that of the proofs of Propositions 11 and 12. First, consider a vertex $v \in I_G$. We have that $\mathcal{A}_{G,\{v\}} = \{V \setminus \{v\}\}$. Now, consider a vertex $v \in S$. In this case, $\mathcal{A}_{G,\{v\}} = \{N_G(v), V \setminus N_G[v]\} = \{N_{\langle S \rangle_G}(v), V \setminus N_G[v]\}$. In consequence, we have that G satisfies $(k, 1)$-adjacency anonymity with

$$\begin{aligned}
k &= \min \left\{ n - 1, \min_{v \in S} \left\{ \min \left\{ \delta_{\langle S \rangle_G}(v), n - \delta_G(v) - 1 \right\} \right\} \right\} \\
&= \min_{v \in S} \left\{ \min \left\{ \delta_{\langle S \rangle_G}(v), n - \delta_G(v) - 1 \right\} \right\} \\
&= \min \left\{ \min_{v \in S} \left\{ \delta_{\langle S \rangle_G}(v), n - \delta_G(v) - 1 \right\} \right\} \\
&= \min \left\{ \min_{v \in S} \left\{ \delta_{\langle S \rangle_G}(v) \right\}, \min_{v \in S} \left\{ n - \delta_G(v) - 1 \right\} \right\} \\
&= \min \left\{ \delta(\langle S \rangle_G), n - \Delta(G) - 1 \right\}
\end{aligned}$$

□

According to Propositions 11, 12 and 13, in order to enforce $(k, 1)$-adjacency anonymity on a graph G, it is necessary to transform it into a graph G' such that the induced subgraph $\langle V(G') \setminus (D_{G'} \cup I_{G'}) \rangle_{G'}$ has minimum degree greater than or equal to k and maximum degree smaller than or equal to $n - k - 1$, where n is the order of G'. Likewise, in order to obtain a $(k, 1)$-adjacency anonymous transformation $(G = (V, E), G' = (V', E'))$, it is necessary to guarantee that every $v \in V$ such that $1 \leq \delta_G(v) < k$ or $|V| - k - 1 < \delta_G(v) \leq |V| - 2$ satisfies $v \notin V'$, or $v \in D_{G'}$, or $v \in I_{G'}$, or $k \leq \delta_{G'}(v) \leq |V'| - k - 1$. Based on these facts, we propose an algorithm that, given a $(k_0, 1)$-adjacency anonymous graph $G = (V, E)$ and an integer k such that $k_0 < k \leq \left\lfloor \frac{|V|-1}{2} \right\rfloor$, efficiently obtains a graph $G' = (V, E')$ such that the pair (G, G') is a $(k, 1)$-adjacency anonymous transformation. The method works by performing a series of edge additions and removals upon G, as outlined in Algorithm 1.

In Algorithm 1, the sets L and H contain the vertices whose degrees are, respectively, smaller and greater than required for the privacy requirement to be satisfied (without being isolated nor dominant vertices). Steps 2, 5, 7, 8, 9, 12, 13, 14, 17, 19, 21, 23, 24, 25, 28, 29, 30 and 33 can be efficiently performed by maintaining the elements of V sorted by their degree and updating the ordering when necessary. The rationale behind the loop in steps 4 to 18 is to first add as many edges as possible between pairs of vertices from L, since every addition of this type increases the degree of two of such vertices. When such additions are no longer possible, then we add edges linking a vertex $u \in L$ and a vertex $v \notin L$ whose degree is as small as possible. The latter condition makes that the degree of vertices from H is only increased if there is no vertex in $V \setminus H$ to which u can be linked. An analogous idea is applied in the loop in steps 20 to 34 to first remove edges joining pairs of vertices from H, then edges joining a vertex from H to other vertex (with the particularity that step 29 takes care

Algorithm 1 Given a $(k_0, 1)$-adjacency anonymous graph $G = (V, E)$ and an integer $k \in \left[k_0 + 1, \left\lfloor \frac{|V|-1}{2} \right\rfloor\right]$, obtain a graph $G' = (V, E')$ such that (G, G') is a $(k, 1)$-adjacency anonymous transformation.

$E' \leftarrow E$
$L \leftarrow \{v \in V : \ 1 \leq \delta_G(v) < k\}$
$H \leftarrow \{v \in V : \ |V| - k - 1 < \delta_G(v) \leq |V| - 2\}$
while $L \neq \emptyset$ **do**
  $X \leftarrow \{x \in L : \ \exists y \in L \setminus \{x\} \text{ s.t. } (x, y) \notin E'\}$
  if $X \neq \emptyset$ **then**
    $u \leftarrow \arg\max_{x \in X}\{\delta_{G'}(x)\}$
    $Y \leftarrow \{y \in L \setminus \{u\} : \ (u, y) \notin E'\}$
    $v \leftarrow \arg\max_{y \in Y}\{\delta_{G'}(y)\}$
    $E' \leftarrow E' \cup \{(u, v)\}$
  else
    $u \leftarrow \arg\max_{x \in L}\{\delta_{G'}(x)\}$
    $Y \leftarrow \{y \in V \setminus L : \ (u, y) \notin E'\}$
    $v \leftarrow \arg\min_{y \in Y}\{\delta_{G'}(y)\}$
    $E' \leftarrow E' \cup \{(u, v)\}$
  end if
  $L \leftarrow \{v \in L : \ 1 \leq \delta_{G'}(v) < k\}$
end while
$H \leftarrow \{v \in H : \ |V| - k - 1 < \delta_{G'}(v) \leq |V| - 2\}$
while $H \neq \emptyset$ **do**
  $X \leftarrow \{x \in H : \ \exists y \in H \setminus \{x\} \text{ s.t. } (x, y) \in E'\}$
  if $X \neq \emptyset$ **then**
    $u \leftarrow \arg\min_{x \in X}\{\delta_{G'}(x)\}$
    $Y \leftarrow \{y \in H \setminus \{u\} : \ (u, y) \in E'\}$
    $v \leftarrow \arg\min_{y \in Y}\{\delta_{G'}(y)\}$
    $E' \leftarrow E' \setminus \{(u, v)\}$
  else
    $u \leftarrow \arg\min_{x \in H}\{\delta_{G'}(x)\}$
    $Y \leftarrow \{y \in V \setminus (H \cup L) : \ (u, y) \in E'\}$
    $v \leftarrow \arg\max_{y \in Y}\{\delta_{G'}(y)\}$
    $E' \leftarrow E' \setminus \{(u, v)\}$
  end if
  $H \leftarrow \{v \in H : \ |V| - k - 1 < \delta_{G'}(v) \leq |V| - 2\}$
end while
return G'

of not making the degree of a vertex from L decrease again), and so on. It is worth noting that in real-life social graphs, which are characterised by very low densities, and for practical values of k, steps 20 to 34 are very unlikely to be executed.

Considering the number of modifications performed by the algorithm, the best scenario is when all edge additions are done according to steps 5 to 10, and all edge removals are done according to steps 21 to 26, as shown in the following results.

Theorem 14. *Let $G = (V, E)$ be a $(k_0, 1)$-adjacency anonymous social graph and let $k \in \left[k_0 + 1, \left\lfloor \frac{|V|-1}{2} \right\rfloor\right]$. The number t of edges added by steps 4 to 18 of Algorithm 1 satisfies*

$$\left\lceil \frac{\sum\limits_{u \in V,\ 1 \leq \delta_G(u) < k} k - \delta_G(u)}{2} \right\rceil \leq t \leq \sum_{u \in V,\ 1 \leq \delta_G(u) < k} k - \delta_G(u) \tag{2}$$

Proof. Let $((u_1, v_1), (u_2, v_2), \ldots, (u_t, v_t))$, with $(u_i, v_i) \in (V \times V) \setminus E$ for $i \in \{1, \ldots, t\}$, be the sequence of edges added to G by steps 4 to 18 of Algorithm 1. Let $E_0 = E$ and $E_i = E_{i-1} \cup \{(u_i, v_i)\}$, for $i \in \{1, \ldots, t\}$. Moreover, for every $i \in \{0, \ldots, t\}$, let $G_i = (V, E_i)$ and $L_i = \{v \in L : 1 \leq \delta_{G_i}(v) < k\}$.

After adding the edge (u_i, v_i), we have that $\delta_{G_i}(u_i) = \delta_{G_{i-1}}(u_i) + 1$ and $\delta_{G_i}(v_i) = \delta_{G_{i-1}}(v_i) + 1$, whereas $\delta_{G_i}(x) = \delta_{G_{i-1}}(x)$ for every $x \in V - \{u_i, v_i\}$.

We define the function

$$missing(G_i) = \sum_{x \in L_i} (k - \delta_{G_i}(x))$$

which specifies by how much the sum of the degrees of vertices from L needs to be increased for (G, G_i) to be a $(k, 1)$-adjacency anonymous transformation. Note that, by the definition of t, we have that $missing(G_t) = 0$. Moreover, $missing(G_0) = \sum\limits_{u \in V,\ 1 \leq \delta_G(u) < k} k - \delta_G(u)$. After adding the edge (u_i, v_i), the following situations are possible:

1. $u_i, v_i \in L_{i-1}$. In this case, since two vertices from L_{i-1} have their degree increased by 1, we have that $missing(G_i) = missing(G_{i-1}) - 2$.

2. $u_i \in L_{i-1}$ and $v_i \notin L_{i-1}$, or *vice versa*. Here, $missing(G_i) = missing(G_{i-1}) - 1$.

With the previous definitions in mind, we will address the proof of the left-hand inequality in Equation 2. To that end, we will assume, for the purpose of contradiction, that

$$t < \left\lceil \frac{\sum_{u \in V,\ 1 \leq \delta_G(u) < k} k - \delta_G(u)}{2} \right\rceil = \left\lceil \frac{missing(G_0)}{2} \right\rceil.$$

If $missing(G_0)$ is even, we have that $t < \frac{missing(G_0)}{2}$. Given that, in the best case scenario, situation 1 above occurs at every iteration of the algorithm, we have

$$\begin{aligned} missing(G_t) &\geq missing(G_0) - 2t \\ &> missing(G_0) - 2 \cdot \tfrac{missing(G_0)}{2} \\ &= 0 \end{aligned}$$

which is a contradiction.

In a similar manner, if $missing(G_0)$ is odd, we have that $t < \frac{missing(G_0)+1}{2}$. Here, in the best case scenario, situation 1 above occurs in every iteration, except one, so

$$\begin{aligned} missing(G_t) &\geq missing(G_0) - 2(t-1) - 1 \\ &> missing(G_0) - 2\left(\tfrac{missing(G_0)+1}{2} - 1\right) - 1 \\ &= 0 \end{aligned}$$

which is also a contradiction. Thus, we can conclude that

$$t \geq \left\lceil \frac{\sum_{u \in V,\ 1 \leq \delta_G(u) < k} k - \delta_G(u)}{2} \right\rceil .$$

The right-hand inequality in Equation 2 is trivial, given that at least one vertex has its degree increased by 1 at every iteration. The proof is thus complete. □

The lower and upper bounds provided in Theorem 14 are tight, as exemplified in Figures 4.a) and 4.b), respectively.

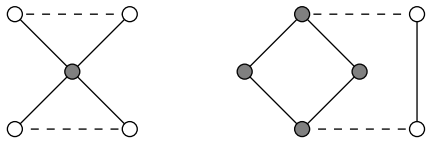

a) b)

Figure 4: Two examples where the number of edges added by steps 4 to 18 of Algorithm 1 (for $k = 2$) reaches the (a) lower and (b) upper bounds of the inequalities in Equation 2. In both cases, dashed lines indicate the edges added by the algorithm.

The next result describes the number t' of edges removed by steps 20 to 34 of Algorithm 1.

Theorem 15. *Let $G = (V, E)$ be a $(k_0, 1)$-adjacency anonymous social graph and let $k \in \left[k_0 + 1, \left\lfloor \frac{|V|-1}{2} \right\rfloor\right]$. Let G_t be the graph obtained from G after executing steps 4 to 18 of Algorithm 1. The number t' of edges removed by steps 20 to 34 of Algorithm 1 satisfies*

$$t' \geq \left\lceil \frac{\sum_{u \in V,\ |V|-k-1 < \delta_G(u) \leq |V|-2} [k - (|V| - \delta_{G_t}(u) - 1)]}{2} \right\rceil \tag{3}$$

and

$$t' \leq \sum_{u \in V,\ |V|-k-1<\delta_G(u)\leq |V|-2} [k - (|V| - \delta_{G_t}(u) - 1)] \tag{4}$$

Proof. We will follow a reasoning analogous to the one applied in the proof of Theorem 14. Let $((u_1, v_1), (u_2, v_2), \ldots, (u_{t'}, v_{t'}))$, with $(u_i, v_i) \in E_t = E(G_t)$ for $i \in \{1, \ldots, t'\}$, be the sequence of edges removed from G_t by steps 20 to 34 of Algorithm 1. Let $E_{t+i} = E_{t+i-1} \setminus \{(u_i, v_i)\}$, for $i \in \{1, \ldots, t'\}$. Moreover, for every $i \in \{1, \ldots, t'\}$, let $G_{t+i} = (V, E_{t+i})$ and $H_i = \{v \in H : |V| - k - 1 < \delta_{G_{t+i}}(v) \leq |V| - 2\}$.

After removing the edge (u_i, v_i), we have that $\delta_{G_{t+i}}(u_i) = \delta_{G_{t+i-1}}(u_i) - 1$ and $\delta_{G_{t+i}}(v_i) = \delta_{G_{t+i-1}}(v_i) - 1$, whereas $\delta_{G_{t+i}}(x) = \delta_{G_{t-i-1}}(x)$ for every $x \in V - \{u_i, v_i\}$.

Now we introduce the function

$$excess(G_{t+i}) = \sum_{x \in H_i} \left[k - (|V| - \delta_{G_{t+i}}(x) - 1)\right].$$

In a manner analogous to the proof of Theorem 14, we have that by definition $excess(G_{t+t'}) = 0$ and $excess(G_t) = \sum_{u \in V,\ |V|-k-1<\delta_G(u)<|V|-2} [k - (|V| - \delta_G(u) - 1)]$. Additionally, after removing the edge (u_i, v_i), the following situations are possible:

1. $u_i, v_i \in H_{i-1}$. In this case, since two vertices from H_{i-1} have their degree decreased by 1, we have that $excess(G_{t+i}) = excess(G_{t+i-1}) - 2$.

2. $u_i \in H_{i-1}$ and $v_i \notin H_{i-1}$, or *vice versa*. Here, $excess(G_{t+i}) = excess(G_{t+i-1}) - 1$.

Now, to address the proof of the inequality in Equation 3, we assume, for the purpose of contradiction, that

$$t' < \left\lceil \frac{\sum_{u \in V,\ |V|-k-1<\delta_G(u)\leq |V|-2} [k - (|V| - \delta_{G_t}(u) - 1)]}{2} \right\rceil = \left\lceil \frac{excess(G_t)}{2} \right\rceil.$$

In consequence, if $excess(G_t)$ is even, we have

$$\begin{aligned} excess(G_{t+t'}) &\geq excess(G_t) - 2t' \\ &> excess(G_t) - 2 \cdot \tfrac{excess(G_t)}{2} \\ &= 0 \end{aligned}$$

which is a contradiction, whereas in the case that $excess(G_t)$ is odd we have

$$\begin{aligned} excess(G_{t+t'}) &\geq excess(G_t) - 2(t' - 1) - 1 \\ &> excess(G_t) - 2\left(\tfrac{excess(G_t)+1}{2} - 1\right) - 1 \\ &= 0 \end{aligned}$$

which is also a contradiction, so we can conclude that Equation 3 holds. As in Theorem 14, the upper bound (Equation 4) is trivial. □

4 $(2,\ell)$-adjacency anonymous transformations

In Algorithm 1, the fact that a vertex v satisfies $v \in L$ means that the equivalence class composed by the vertices having adjacency representation (1) with respect to the set $\{v\}$ in the original graph G is not empty and its cardinality is smaller than k. Likewise, the fact that $v \in H$ means that the equivalence class composed by the vertices having adjacency representation (2) with respect to the set $\{v\}$ in G is not empty and its cardinality is smaller than k. To a limited extent, a strategy similar to the one applied in Algorithm 1 can be used to obtain (k,ℓ)-adjacency anonymous transformations with $\ell > 1$. For example, for $\ell = 2$, in addition to the sets $L, H \subset V(G)$, we would consider the sets $P_{11}, P_{12}, P_{21}, P_{22} \subseteq V(G) \times V(G)$, where $(u,v) \in P_{ij}$, $i,j \in \{1,2\}$, means that the the equivalence class composed by the vertices having adjacency representation (i,j) with respect to (u,v) in G is not empty and its cardinality is smaller than k. Thus, the algorithm would work by executing the necessary edge set editions to increase the cardinalities of these equivalence classes or, alternatively, to empty them. However, it is impractical to use this philosophy in the general case, as it entails designing a different, highly casuistic algorithm for every different value of ℓ.

For the general case, we have devised a greedy edge-addition-based method that, for small values of ℓ, allows to obtain $(2,\ell)$-adjacency anonymous transformations. Given a graph $G = (V,E)$, the method starts by computing all 1-adjacency antiresolving sets of G of cardinality at most ℓ. Then, edges are iteratively added until obtaining a graph $G' = (V, E \cup E')$ such that (G,G') is a $(2,\ell)$-adjacency anonymous transformation. The critical aspect of the proposed framework is how to determine an appropriate order for adding edges.

In order to describe the proposed method, we will first introduce the following results, which characterise the sets of edges whose addition to a graph G may modify the set of 1-adjacency antiresolving sets. In what follows, we will use the notation $\mathcal{S}_{G,\ell}$ for the set of 1-adjacency antiresolving sets of a graph G having cardinality smaller than or equal to ℓ.

Remark 16. *Let $G = (V,E)$ be a social graph, $u,v \in V$ a pair of vertices of G such that $(u,v) \notin E$, and $G' = (V, E \cup \{(u,v)\})$. If, for every $S \in \mathcal{S}_{G,\ell}$, either $S \cap \{u,v\} = \{u,v\}$ or $S \cap \{u,v\} = \emptyset$, then $\mathcal{S}_{G,\ell} \setminus \mathcal{S}_{G',\ell} = \mathcal{S}_{G,\ell}$.*

Proof. The result follows directly from the fact that, for every $S \in \mathcal{S}_{G,\ell}$ and every $x \in V \setminus S$, we have that $a_G(x \mid S) = a_{G'}(x \mid S)$, so $S \in \mathcal{S}_{G',\ell}$. □

Algorithm 2 describes the edge-addition method. First, we use Remark 16 to discard candidate vertex pairs (u,v) that are known not to cause any 1-adjacency antiresolving set of the current graph $G = (V,E)$ to become a k-adjacency antiresolving set of $G' = (V, E \cup \{(u,v)\})$ with $k > 1$. Then, every remaining candidate pair (u,v) is scored as follows:

$$score((u,v)) = |\{S : \ S \in \mathcal{S}_{G,\ell}, u \in S, [v]_S^G = \{v\}\}|.$$

where $[v]_S^G$ represents the equivalence class of v in $\mathcal{A}_{G,S}$. In other words, we consider the number of times the candidate pair would modify the fingerprint of a uniquely identifiable vertex with respect to a 1-adjacency antiresolving set of G. The intuition behind this heuristics is that the larger the number of times that the pair (u, v) is found in this situation, the larger the likelihood that adding the edge (u, v) will result in making some vertex set stop being 1-adjacency antiresolving. At every iteration, the current perturbed graph $G' = (V, E')$ is transformed into the graph $G'' = (V, E' \cup \{(u, v)\})$, where (u, v) is the best-scored candidate addition satisfying $(\mathcal{S}_{G'',\ell} \setminus \mathcal{S}_{G',\ell}) \cap \mathcal{S}_{G,\ell} = \emptyset$.

Algorithm 2 Given a graph $G = (V, E)$ and a positive integer $\ell \geq 2$, obtain a graph $G' = (V, E')$ such that (G, G') is a $(2, \ell)$-anonymous transformation.

1: $C \leftarrow \emptyset$
2: Compute $\mathcal{S}_{G,\ell}$
3: $C \leftarrow \emptyset$
4: **for** $(u, v) \in (V \times V) \setminus E$ **do**
5: **for** $S \in \mathcal{S}_{G,\ell}$ **do**
6: **if** $|S \cap (u, v)| = 1$ **then**
7: $C \leftarrow C \cup \{(u, v)\}$
8: **break**
9: **end if**
10: **end for**
11: **end for**
12: Decrementally sort C by $score((u, v)) = |\{S : S \in \mathcal{S}_{G,\ell}, u \in S, [v]_S^G = \{v\}\}|$
13: $t \leftarrow 0$
14: $E_0 \leftarrow E$
15: **while** $\mathcal{S}_{G,\ell} \setminus \mathcal{S}_{G_t,\ell} \neq \emptyset$ **do**
16: $(u, v) = \arg\max_{score(u,v)}\{(u, v) \in C : (\mathcal{S}_{(V,E_t \cup \{(u,v)\}),\ell} \setminus \mathcal{S}_{G_t,\ell}) \cap \mathcal{S}_{G,\ell} = \emptyset\}$
17: $E_{t+1} \leftarrow E_t \cup \{(u, v)\}$
18: $C \leftarrow C \setminus \{(u, v)\}$
19: $t \leftarrow t + 1$
20: **end while**
21: **return** G_t

The asymptotic time complexity of Algorithm 2 is dominated by the computation and traversals of $\mathcal{S}_{G,\ell}$, which is $O(2^n)$ in the general case. However, for small values of ℓ, these computations can be done in $O(n^\ell)$ time. As we discussed before, active adversaries can only insert a limited amount of sybil nodes in the network without being detected, so the capacity of protecting the graph against such adversaries results in an important privacy increase. As an aid to speed-up the algorithm, the following result shows how the number of verifications to perform in evaluating the condition at step 16 can be largely reduced.

Theorem 17. *Let $G = (V, E)$ be a social graph, $u, v \in V$ a pair of vertices of G such that $(u, v) \notin E$, and $G' = (V, E \cup \{(u, v)\})$. Let $S \subseteq V$ such that it is a k-adjacency*

antiresolving set of G, with $k \geq 2$, and a 1-adjacency antiresolving set of G'. If there exists $w \in S$ such that $d_G(u,w) > 2$ and $d_G(v,w) > 2$, then for every $x \in V \setminus S$ such that $[x]_S^{G'} = \{x\}$ it holds that $[x]_{S\setminus\{w\}}^{G'} = \{x\}$ and $|[x]_{S\setminus\{w\}}^{G}| > 1$.

Proof. Consider a graph $G = (V, E)$, a pair of vertices $u, v \in V$ such that $(u, v) \notin E$ and a set $S \subseteq V$ satisfying the premises of Theorem 17. Also consider a vertex $w \in S$ such that $d_G(u, w) > 2$ and $d_G(v, w) > 2$ and a vertex $x \in V \setminus S$ such that $[x]_S^{G'} = \{x\}$. By the definition of S we have that $|[x]_S^G| > 1$, so there exists $y \in V \setminus (S \cup \{x\})$ such that $a_G(x|\ S) = a_G(y|\ S)$ and $a_{G'}(x|\ S) \neq a_{G'}(y|\ S)$. Suppose, for the purpose of contradiction, that $a_{G'}(w, x) \neq a_{G'}(w, y)$. Then, since $w \neq u$ and $w \neq v$, we have that $a_G(w, x) \neq a_G(w, y)$, which contradicts the fact that $a_G(x|\ S) = a_G(y|\ S)$. Therefore, we have that $a_{G'}(w, x) = a_{G'}(w, y)$, which implies that $a_G(x|\ S\setminus\{w\}) = a_G(y|\ S\setminus\{w\})$ and $a_{G'}(x|\ S \setminus \{w\}) \neq a_{G'}(y|\ S \setminus \{w\})$. Since the only difference between G and G' is the addition of the edge (u, v), we conclude that either $x = u$ and $v \in S$, or *vice versa*. For the remainder of this proof, we will assume $x = u$ and $v \in S$ without loss of generality.

We will now proceed by reduction to absurdity. To that end, we will assume that $|[x]_{S\setminus\{w\}}^{G'}| > 1$ or $[x]_{S\setminus\{w\}}^{G} = \{x\}$. It is simple to see that $[x]_{S\setminus\{w\}}^{G} = \{x\}$ contradicts the fact that S is a k-adjacency antiresolving set of G with $k \geq 2$. Therefore, in what follows we will focus on the assumption that $|[x]_{S\setminus\{w\}}^{G'}| > 1$. In this case, there exists $y \in V \setminus (S \cup \{x\})$ such that $a_{G'}(x|\ S \setminus \{w\}) = a_{G'}(y|\ S \setminus \{w\})$. Since $[x]_S^{G'} = \{x\}$, we have that $a_{G'}(w, x) \neq a_{G'}(w, y)$. Moreover, since $(w, x) \neq (u, v)$ and $(w, y) \neq (u, v)$, we have that $a_G(w, x) = a_{G'}(w, x)$ and $a_G(w, y) = a_{G'}(w, y)$. Thus, it holds that $a_{G'}(w, x) \neq a_{G'}(w, y) \implies a_G(w, x) \neq a_G(w, y)$, which entails

$$(w \sim_G x \wedge w \nsim_G y) \vee (w \nsim_G x \wedge w \sim_G y)$$

Since $x = u$, we have that $w \nsim_G x$, because $d_G(u, w) > 2$, so $w \sim_G y$. Moreover, since $u \sim_{G'} v$ and $a_{G'}(v, x) = a_{G'}(v, y)$, then $v \sim_G y$, which implies $d_G(w, v) = 2$, again a contradiction. This concludes the proof. □

According to Theorem 17, when verifying if the addition of an edge (u, v) causes some k-adjacency antiresolving set (with $k \geq 2$) of G_t to become a 1-adjacency antiresolving set of $G'_t = (V(G_t), E(G_t) \cup \{(u, v)\})$, it suffices to analyse those sets $S \in \mathcal{S}_{G_t,\ell}$ such that some $w \in S$ satisfies $d_{G_t}(u, w) \leq 2$ or $d_{G_t}(v, w) \leq 2$.

5 Concluding remarks

In this paper, we have re-visited the notion of (k, ℓ)-anonymity, which quantifies the privacy level of a social graph in the presence of active adversaries. Firstly, we have introduced the notion of (k, ℓ)-anonymous transformations, which allow to reduce the amount of perturbation needed to protect a social graph from an active attack. Secondly, we have critically assessed the assumptions posed by (k, ℓ)-anonymity on the adversary capabilities. Judging that it is unrealistic to assume that an adversary will be able to control all distances between a set of sybil nodes and every other

vertex of the social graph, we introduced a new privacy property: (k, ℓ)-adjacency anonymity, which accounts for adversaries who control the connection patterns with the neighbours of the sybil nodes. Finally, combining the two previous ideas, we have introduced (k, ℓ)-adjacency anonymous transformations, which are able to protect a social graph from active adversaries levaraging up to ℓ sybil nodes and constructing fingerprints based on the connection patterns between victims and sybil nodes. We proposed two algorithms: one for obtaining $(k, 1)$-adjacency anonymous transformations for arbitrary values of k, and another for obtaining $(2, \ell)$-adjacency anonymous transformations for small values of ℓ. The first algorithm is efficient and the number of changes introduced in the graph is bounded. On the other hand, there is still room for improvement in the second method, especially concerning the order in which graph perturbations are applied. We are currently using a greedy heuristic to guide the edge-addition process. We will evaluate the convenience of this heuristic, and explore the use of meta-heuristics such as genetic algorithms and ant-colony optimisation.

References

[1] L. Backstrom, C. Dwork, J. Kleinberg, Wherefore art thou r3579x?: anonymized social networks, hidden patterns, and structural steganography, in: Proceedings of the 16th international conference on World Wide Web, WWW '07, ACM, New York, NY, USA, 2007, pp. 181–190.
URL http://doi.acm.org/10.1145/1242572.1242598

[2] F. Harary, R. A. Melter, On the metric dimension of a graph, Ars Combinatoria 2 (1976) 191–1995.

[3] M. Jannesari, B. Omoomi, The metric dimension of the lexicographic product of graphs, Discrete Mathematics 312 (22) (2012) 3349–3356.
URL http://doi.org/10.1016/j.disc.2012.07.025

[4] S. Mauw, Y. Ramírez-Cruz, R. Trujillo-Rasua, An edge addition approach to counteract active attacks in social graphs, *Submitted manuscript*, 2016.

[5] S. Mauw, R. Trujillo-Rasua, B. Xuan, Counteracting active attacks in social network graphs, in: Proceedings of DBSec'16, vol. 9766 of Lecture Notes in Computer Science, 2016, pp. 233–248.
URL https://doi.org/10.1007/978-3-319-41483-6_17

[6] A. Narayanan, V. Shmatikov, De-anonymizing social networks, in: Procs. of the 30th IEEE Symposium on Security and Privacy, 2009, pp. 173–187.
URL http://dx.doi.org/10.1109/SP.2009.22

[7] W. Peng, F. Li, X. Zou, J. Wu, Seed and grow: An attack against anonymized social networks, in: Proc. of SECON 2012, 2012, pp. 587–595.

[8] W. Peng, F. Li, X. Zou, J. Wu, A two-stage deanonymization attack against anonymized social networks, IEEE Transactions on Computers 63 (2) (2014) 290–303.

[9] A. Sebő, E. Tannier, On metric generators of graphs, Mathematics of Operations Research 29 (2) (2004) 383–393.
URL http://dx.doi.org/10.1287/moor.1030.0070

[10] P. J. Slater, Leaves of trees, Congressus Numerantium 14 (1975) 549–559.

[11] R. Trujillo-Rasua, I. G. Yero, k-metric antidimension: A privacy measure for social graphs, Inf. Sci. 328 (2016) 403–417.
URL http://dx.doi.org/10.1016/j.ins.2015.08.048